

ThinkSystem NVIDIA RTX PRO 6000 Blackwell Server Edition PCIe Gen5 GPU

Product Guide

The ThinkSystem NVIDIA RTX PRO 6000 Blackwell Server Edition PCIe Gen5 GPU is a powerful Blackwell data center platform for AI and visual computing, delivering breakthrough performance to accelerate a wide array of enterprise data center workloads. Enterprises can configure RTX PRO 6000 Blackwell GPU in high-density server platforms to deliver unmatched levels of compute power, memory capacity, and throughput to power mission-critical AI-enabled applications and accelerate use cases across industries—from healthcare, manufacturing, and geoscience to retail, media, and live broadcast.




Figure 1. ThinkSystem NVIDIA RTX PRO 6000 Blackwell Server Edition PCIe Gen5 GPU

Did you know?

As part of the NVIDIA OVX server platform, RTX PRO 6000 Blackwell GPU delivers the highest level of graphics, ray tracing, and simulation performance for NVIDIA Omniverse. With 96 GB of GDDR7 memory, even the most intense graphics applications run at the highest level of performance.

Part number information

The following table shows the part numbers for the RTX PRO 6000 Blackwell GPU.

Table 1. Ordering information

Part number	Feature code	Description	NVIDIA part number	Controlled GPU status
4X67B09287	CBK8	ThinkSystem NVIDIA RTX PRO 6000 Blackwell Server Edition 96GB PCIe Gen5 Passive GPU	900-2G153-2700-030	Controlled

The RTX PRO 6000 Blackwell GPU is Controlled which means the GPU is not offered in certain markets, as determined by the US Government.

The PCIe option part numbers includes the following:

- One RTX PRO 6000 Blackwell GPU with full-height (3U) adapter bracket attached
- Documentation

Features

Based on the NVIDIA Blackwell architecture, the RTX PRO 6000 Blackwell GPU is a full-height, full-length (FHFL), dual-slot, 10.5-inch PCIe Gen5 graphics solution. The card is passively cooled and is capable of 600 W maximum board power. Equipped with 96 GB of ultra-fast GDDR7 memory, the RTX PRO 6000 Blackwell GPU provides unparalleled performance and flexibility to accelerate a broad range of use cases—from multimodal AI inference, physical AI, and scientific computing to rendering, 3D graphics, video, and more.

Suitable workloads for the RTX PRO 6000 Blackwell GPU include the following:

- LLM Fine-Tuning, Inference, and Retrieval-Augmented Generation (RAG) workloads with fewer than 70B parameters
- Visual computing workloads such as graphics, video, architecture, and product design
- Workloads such as data analytics, FP32 High Performance Computing

Key features of the RTX PRO 6000 Blackwell GPU:

- Fifth-Generation Tensor cores, FP4 precision
- Fourth-Generation ray-tracing cores
- 96 GB of GDDR7 memory
- 1597 GB/s of memory bandwidth
- 9th Generation NVENC and 6th Generation NVDEC with 4:2:2 support
- PCIe Gen 5 x16 host interface
- Multi-instance GPU (MIG) support
- Passive thermal design

The ThinkSystem NVIDIA RTX PRO 6000 Blackwell Server Edition PCIe Gen5 GPU offers the following features:

- Fifth-Generation Tensor Cores:
Deliver up to 5x the performance of the previous generation for LLM inference with support for FP4 precision for faster AI model processing times and reduced memory usage.
- Fourth-Generation Ray-Tracing Cores:
Double the ray-triangle intersection rate of the previous generation to create photo-realistic, physically accurate scenes and immersive 3D designs with RTX™ Mega Geometry, which enables up to 100x more ray-traced triangles.

- Next-Gen Video Engines:

Accelerate multimodal AI inference, enhance video production, and power streaming workflows with real-time AI processing. The RTX PRO 6000 Blackwell features four 9th Generation NVENC and four 6th Generation NVDEC engines with support for 4:2:2 encoding and decoding.

- GDDR7 Memory:

New and improved GDDR7 memory significantly boosts bandwidth and capacity, empowering your applications to run faster and work with larger, more complex datasets. With 96 GB of GPU memory and 1.6 TB/s bandwidth, enterprises can deliver low-latency inference with larger models and render massive 3D scenes.

- NVIDIA Deep Learning Super Sampling (DLSS) 4:

DLSS is a revolutionary suite of neural rendering technologies that uses AI to boost FPS, reduce latency, and improve image quality. Multi-Frame Generation ensures ultra-smooth frame pacing for lifelike simulations. Experience up to 3x faster frame rates and stunning image quality in supported game engines and 3D rendering applications for smoother, more responsive performance.

- PCIe Gen 5:

Support for PCIe Gen 5 provides double the bandwidth of PCIe Gen 4, improving data-transfer speeds from CPU memory and unlocking faster performance for data-intensive tasks like AI, data science, and 3D modeling.

- Universal MIG:

Divide a single RTX PRO 6000 Blackwell into up to four fully isolated instances, each with dedicated resources, allowing for concurrent execution of graphics and AI workloads, optimized GPU utilization, and secure isolation of different applications or users.

Technical specifications

The following table lists the specifications of the RTX PRO 6000 Blackwell GPU.

Table 2. RTX PRO 6000 Blackwell GPU specifications

Feature	Specification
GPU Memory	96 GB GDDR7
Memory Interface	512-bit
Memory Bandwidth	Up to 1597 GB/s
ECC	Yes
NVIDIA CUDA Cores	24,064
NVIDIA Tensor Cores	752 (5th Generation)
NVIDIA RT Cores	188 (4th Generation)
FP32 Single-Precision Performance	117 TFLOPS (peak)
FP4 AI performance	3.7 PFLOPS (peak)
RT Core performance	354.5 TFLOPS (peak)
Host Interface	PCIe 5.0 x16
NVLink support	No
Power Consumption	600 W
Multi-Instance GPU (MIG)	Up to 4 MIGs @ 24GB
Thermal Solution	Passive cooling
Form Factor	4.4" H x 10.5" L, Dual Slot, Full Height, Full Length
Display Connectors	4x DisplayPort 2.1b (disabled by default**)
Maximum simultaneous displays	4x 4K displays at 165 Hz using DP 2x 4K displays 360 Hz with DSC using DP 2x 8K displays at 100 Hz with DSC using DP
Encode / Decode Engines	4x NVENC encode, 4x NVDEC decode, 4x JPEG
vGPU software support	NVIDIA vPC/vApps, NVIDIA RTX Virtual Workstation (vWS)
Graphics APIs	DirectX 12, Shader Model 6.6, OpenGL 4.6, Vulkan 1.3

** To enable the DisplayPort ports, see <https://developer.nvidia.com/displaymodeselector>

Server support

The following tables list the ThinkSystem servers that are compatible.

Table 3. Server support (Part 1 of 5)

Part Number	Description	AMD V3				2S Intel V3/V4				Multi Node V3	1S V3							
		SR635 V3 (7D9H / 7D9G)	SR655 V3 (7D9F / 7D9E)	SR645 V3 (7D9D / 7D9C)	SR665 V3 (7D9B / 7D9A)	ST650 V3 (7D7B / 7D7A)	SR630 V3 (7D72 / 7D73)	SR650 V3 (7D75 / 7D76)	SR630 V4 (7DG8 / 7DG9)	SR650 V4 (7DGC / 7DGD)	SR650a V4 (7DGC / 7DGD)	SD535 V3 (7DD8 / 7DD1)	SD530 V3 (7DDA / 7DD3)	SD550 V3 (7DD9 / 7DD2)	ST45 V3 (7DH4 / 7DH5)	ST50 V3 (7DF4 / 7DF3)	ST250 V3 (7DCF / 7DCE)	SR250 V3 (7DCM / 7DCL)
4X67B09287	ThinkSystem NVIDIA RTX PRO 6000 Blackwell Server Edition 96GB PCIe Gen5 Passive GPU	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N

Table 4. Server support (Part 2 of 5)

Part Number	Description	4S 8S Intel V3/V4		GPU Rich						Edge									
		SR850 V3 (7D97 / 7D96)	SR860 V3 (7D94 / 7D93)	SR950 V3 (7DC5 / 7DC4)	SR850 V4 (7DJT / 7DJS)	SR860 V4 (7DJQ / 7DJN)	SR670 V2 (7Z22 / 7Z23)	SR675 V3 (7D9Q / 7D9R)	SR680a V3 (7DHE)	SR680a V3 B200 (7DM9)	SR685a V3 (7DHC)	SR780a V3 (7DJ5)	SR680a V4 (7DMK)	SE100 (7DGR)	SE350 (7Z46 / 7D1X)	SE350 V2 (7DA9)	SE360 V2 (7DAM)	SE450 (7D8T)	SE455 V3 (7DBY)
4X67B09287	ThinkSystem NVIDIA RTX PRO 6000 Blackwell Server Edition 96GB PCIe Gen5 Passive GPU	N	N	N	N	N	8	N	N	N	N	N	N	N	N	N	N	N	N

Table 5. Server support (Part 3 of 5)

Part Number	Description	Super Computing							1S Intel V2		2S Intel V2		AMD V1					
		SC750 V4 (7DDJ)	SC777 V4 (7DKA)	SD665 V3 (7D9P)	SD665-N V3 (7DAZ)	SD650 V3 (7D7M)	SD650-I V3 (7D7L)	SD650-N V3 (7D7N)	ST50 V2 (7D8K / 7D8J)	ST250 V2 (7D8G / 7D8F)	SR250 V2 (7D7R / 7D7Q)	ST650 V2 (7Z75 / 7Z74)	SR630 V2 (7Z70 / 7Z71)	SR650 V2 (7Z72 / 7Z73)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR645 (7D2Y / 7D2X)	SR665 (7D2W / 7D2V)
4X67B09287	ThinkSystem NVIDIA RTX PRO 6000 Blackwell Server Edition 96GB PCIe Gen5 Passive GPU	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N

Table 6. Server support (Part 4 of 5)

Part Number	Description	Dense V2				4S V2	8S	4S V1		1S Intel V1				
		SD630 V2 (7D1K)	SD650 V2 (7D1M)	SD650-N V2 (7D1N)	SN550 V2 (7Z69)	SR850 V2 (7D31 / 7D32)	SR860 V2 (7Z59 / 7Z60)	SR950 (7X11 / 7X12)	SR850 (7X18 / 7X19)	SR850P (7D2F / 2D2G)	SR860 (7X69 / 7X70)	ST50 (7Y48 / 7Y50)	ST250 (7Y45 / 7Y46)	SR150 (7Y54)
4X67B09287	ThinkSystem NVIDIA RTX PRO 6000 Blackwell Server Edition 96GB PCIe Gen5 Passive GPU	N	N	N	N	N	N	N	N	N	N	N	N	N

Table 7. Server support (Part 5 of 5)

Part Number	Description	2S Intel V1								Dense V1			
		ST550 (7X09 / 7X10)	SR530 (7X07 / 7X08)	SR550 (7X03 / 7X04)	SR570 (7Y02 / 7Y03)	SR590 (7X98 / 7X99)	SR630 (7X01 / 7X02)	SR650 (7X05 / 7X06)	SR670 (7Y36 / 7Y37)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)
4X67B09287	ThinkSystem NVIDIA RTX PRO 6000 Blackwell Server Edition 96GB PCIe Gen5 Passive GPU	N	N	N	N	N	N	N	N	N	N	N	N

NVIDIA GPU software

This section lists the NVIDIA software that is available from Lenovo.

- [NVIDIA vGPU Software \(vApps, vPC, RTX vWS\)](#)
- [NVIDIA Omniverse Software \(OVE\)](#)
- [NVIDIA AI Enterprise Software](#)
- [NVIDIA Run:ai](#)
- [NVIDIA HPC Compiler Software](#)

NVIDIA vGPU Software (vApps, vPC, RTX vWS)

Lenovo offers the following virtualization software for NVIDIA GPUs:

- **Virtual Applications (vApps)**

For organizations deploying Citrix XenApp, VMware Horizon RDSH or other RDSH solutions. Designed to deliver PC Windows applications at full performance. NVIDIA Virtual Applications allows users to access any Windows application at full performance on any device, anywhere. This edition is suited for users who would like to virtualize applications using XenApp or other RDSH solutions. Windows Server hosted RDSH desktops are also supported by vApps.

- **Virtual PC (vPC)**

This product is ideal for users who want a virtual desktop but need great user experience leveraging PC Windows® applications, browsers and high-definition video. NVIDIA Virtual PC delivers a native experience to users in a virtual environment, allowing them to run all their PC applications at full performance.

- **NVIDIA RTX Virtual Workstation (RTX vWS)**

NVIDIA RTX vWS is the only virtual workstation that supports NVIDIA RTX technology, bringing advanced features like ray tracing, AI-denoising, and Deep Learning Super Sampling (DLSS) to a virtual environment. Supporting the latest generation of NVIDIA GPUs unlocks the best performance possible, so designers and engineers can create their best work faster. IT can virtualize any application from the data center with an experience that is indistinguishable from a physical workstation — enabling workstation performance from any device.

The following license types are offered:

- **Perpetual license**

A non-expiring, permanent software license that can be used on a perpetual basis without the need to renew. For each perpetual license, customers are also required to purchase a 5-year SUMS support contract. Without this contract, the perpetual license cannot be ordered.

- **Annual subscription**

A software license that is active for a fixed period as defined by the terms of the subscription license, typically yearly. The subscription includes Support, Upgrade and Maintenance (SUMS) for the duration of the license term.

- **Concurrent User (CCU)**

A method of counting licenses based on active user VMs. If the VM is active and the NVIDIA vGPU software is running, then this counts as one CCU. A vGPU CCU is independent of the connection to the VM.

The following table lists the ordering part numbers and feature codes.

Table 8. NVIDIA vGPU Software

Part number	Feature code 7S02CTO1WW	NVIDIA part number	Description
NVIDIA vApps			
7S020004WW	B1MQ	711-VAP002+P3CMI12	NVIDIA vApps Subscription License 1 Year, 1 CCU

Part number	Feature code 7S02CTO1WW	NVIDIA part number	Description
7S020005WW	B1MR	711-VAP002+P3CMI36	NVIDIA vApps Subscription License 3 Years, 1 CCU
7S02003DWW	S832	711-VAP002+P3CMI48	NVIDIA vApps Subscription License 4 Years, 1 CCU
7S02003EWW	S833	711-VAP002+P3CMI60	NVIDIA vApps Subscription License 5 Years, 1 CCU
7S020046WW	SDHB	711-VAP001+P3CMI00	NVIDIA vApps Perpetual License, 1 CCU
7S020003WW	B1MP	711-VAP001+P3CMI00	NVIDIA vApps SUMS ONLY 5Yr, 1 CCU (required for perpetual license)
NVIDIA vPC			
7S02000Aww	B1MW	711-VPC022+P3CMI12	NVIDIA vPC Subscription License 1 Year, 1 CCU
7S02000BWW	B1MX	711-VPC022+P3CMI36	NVIDIA vPC Subscription License 3 Years, 1 CCU
7S02003FWW	S834	711-VPC022+P3CMI48	NVIDIA vPC Subscription License 4 Years, 1 CCU
7S02003GWW	S835	711-VPC022+P3CMI60	NVIDIA vPC Subscription License 5 Years, 1 CCU
7S020047WW	SDHC	711-VPC021+P3CMI00	NVIDIA vPC Perpetual License, 1 CCU
7S020009WW	B1MV	711-VPC021+P3CMI00	NVIDIA vPC SUMS 5Yr ONLY, 1 CCU (required for perpetual license)
NVIDIA RTX vWS			
7S02000GWW	B1N2	711-DWS022+P3CMI12	NVIDIA RTX vWS Subsc Lic 1Yr 1 CCU
7S02000HWW	B1N3	711-DWS022+P3CMI36	NVIDIA RTX vWS Subscription License 3 Years, 1 CCU
7S02000XWW	S6YJ	711-DWS022+P3CMI48	NVIDIA RTX vWS Subscription License 4 Years, 1 CCU
7S02000YWW	S6YK	711-DWS022+P3CMI60	NVIDIA RTX vWS Subscription License 5 Years, 1 CCU
7S02000MWW	B1N7	711-DWS022+P3EDI12	NVIDIA RTX vWS EDU Subscription License 1 Year, 1 CCU
7S02000NWW	B1N8	711-DWS022+P3EDI36	NVIDIA RTX vWS EDU Subscription License 3 Years, 1 CCU
7S02003BWW	S830	711-DWS022+P3EDI48	NVIDIA RTX vWS EDU Subscription License 4 Years, 1 CCU
7S02003CWW	S831	711-DWS022+P3EDI60	NVIDIA RTX vWS EDU Subscription License 5 Years, 1 CCU
7S020048WW	SDHD	711-DWS021+P3CMI00	NVIDIA RTX vWS Perpetual License, 1 CCU
7S02000FWW	B1N1	711-DWS021+P3CMI00	NVIDIA RTX vWS SUMS ONLY 5Yr, 1 CCU (required for perpetual license)
7S020049WW	SDHE	711-DWS021+P3EDI00	NVIDIA RTX vWS EDU Perpetual License, 1 CCU
7S02000LWW	B1N6	711-DWS021+P3EDI00	NVIDIA RTX vWS EDU SUMS ONLY 5Y, 1CCU (required for perpetual license)
NVIDIA RTX vWS Support & Services			
7S020015WW	S6YS	712-DWSA24+P3CMI12	24X7 Support Services for NVIDIA RTX vWS Production SUMS, 1CCU, 1 Year
7S02005CWW	SDZB	712-DWSA24+P3CMI60	24X7 Support Services for NVIDIA RTX vWS Production SUMS 1CCU 5 Years
7S020016WW	S6YT	712-DWSA24+P3EDI12	24X7 Support Services for NVIDIA RTX vWS Production SUMS, 1CCU, EDU, 1 Year

Part number	Feature code 7S02CTO1WW	NVIDIA part number	Description
7S02005DWW	SDZC	712-DWSA24+P3EDI60	24X7 Support Services for NVIDIA RTX vWS Production SUMS 1CCU EDU 5 Years
7S02005EWW	SDZD	712-DWSB24+P3CMI12	24X7 Support Services for NVIDIA RTX vWS SUMS 4 CCU 1 Year
7S020017WW	S6YU	712-DWSD24+P3CMI12	24X7 Support Services for NVIDIA RTX vWS Subscription License, 1CCU, 1 Year

NVIDIA Omniverse Software (OVE)

NVIDIA Omniverse™ Enterprise is an end-to-end collaboration and simulation platform that fundamentally transforms complex design workflows, creating a more harmonious environment for creative teams.

NVIDIA and Lenovo offer a robust, scalable solution for deploying Omniverse Enterprise, accommodating a wide range of professional needs. This document details the critical components, deployment options, and support available, ensuring an efficient and effective Omniverse experience.

Deployment options cater to varying team sizes and workloads. Using Lenovo NVIDIA-Certified Systems™ and Lenovo OVX nodes which are meticulously designed to manage scale and complexity, ensures optimal performance for Omniverse tasks.

Deployment options include:

- Workstations: NVIDIA-Certified Workstations with RTX 6000 Ada GPUs for desktop environments.
- Data Center Solutions: Deployment with Lenovo OVX nodes or NVIDIA-Certified Servers equipped with L40, or L40S GPUs for centralized, high-capacity needs.

NVIDIA Omniverse Enterprise includes the following components and features:

- Platform Components: Kit, Connect, Nucleus, Simulation, RTX Renderer.
- Foundation Applications: USD Composer, USD Presenter.
- Omniverse Extensions: Connect Sample & SDK.
- Integrated Development Environment (IDE)
- Nucleus Configuration: Workstation, Enterprise Nucleus Server (supports up to 8 editors per scene); Self-Service Public Cloud Hosting using Containers.
- Omniverse Farm: Supports batch workloads up to 8 GPUs.
- Enterprise Services: Authentication (SSO/SSL), Navigator Microservice, Large File Transfer, User Accounts SAML/Account Directory.
- User Interface: Workstation & IT Managed Launcher.
- Support: NVIDIA Enterprise Support.
- Deployment Scenarios: Desktop to Data Center: Workstation deployment for building and designing, with options for physical or virtual desktops. For batch tasks, rendering, and SDG workloads that require headless compute, Lenovo OVX nodes are recommended.

The following part numbers are for a subscription license which is active for a fixed period as noted in the description. The license is for a named user which means the license is for named authorized users who may not re-assign or share the license with any other person.

Table 9. NVIDIA Omniverse Software (OVE)

Part number	Feature 7S02CTO1WW	NVIDIA part number	Description
7S02003ZWW	SCX0	721-OV7006+P3CMI12	NVIDIA Omniverse Enterprise Subscription per GPU, 1 Year
7S020042WW	SCX3	721-OV7006+P3CMI36	NVIDIA Omniverse Enterprise Subscription per GPU, 3 Years
7S020044WW	SD5T	721-OV7006+P3CMI60	NVIDIA Omniverse Enterprise Subscription per GPU, 5 Year
7S020041WW	SCX2	721-OV7006+P3INI12	NVIDIA Omniverse Enterprise Subscription per GPU, INC, 1 Year
7S020040WW	SCX1	721-OV7006+P3EDI12	NVIDIA Omniverse Enterprise Subscription per GPU, EDU, 1 Year
7S020043WW	SCX4	721-OV7006+P3EDI36	NVIDIA Omniverse Enterprise Subscription per GPU, EDU, 3 Years
7S020045WW	SD5U	721-OV7006+P3EDI60	NVIDIA Omniverse Enterprise Subscription per GPU EDU, 5 Year

NVIDIA AI Enterprise Software

Lenovo offers the NVIDIA AI Enterprise (NVAIE) cloud-native enterprise software. NVIDIA AI Enterprise is an end-to-end, cloud-native suite of AI and data analytics software, optimized, certified, and supported by NVIDIA to run on VMware vSphere and bare-metal with NVIDIA-Certified Systems™. It includes key enabling technologies from NVIDIA for rapid deployment, management, and scaling of AI workloads in the modern hybrid cloud.

NVIDIA AI Enterprise is licensed on a per-GPU basis. NVIDIA AI Enterprise products can be purchased as either a perpetual license with support services, or as an annual or multi-year subscription.

- The perpetual license provides the right to use the NVIDIA AI Enterprise software indefinitely, with no expiration. NVIDIA AI Enterprise with perpetual licenses must be purchased in conjunction with one-year, three-year, or five-year support services. A one-year support service is also available for renewals.
- The subscription offerings are an affordable option to allow IT departments to better manage the flexibility of license volumes. NVIDIA AI Enterprise software products with subscription includes support services for the duration of the software's subscription license

The features of NVIDIA AI Enterprise Software are listed in the following table.

Table 10. Features of NVIDIA AI Enterprise Software (NVAIE)

Features	Supported in NVIDIA AI Enterprise
Per GPU Licensing	Yes
Compute Virtualization	Supported
Windows Guest OS Support	No support
Linux Guest OS Support	Supported
Maximum Displays	1
Maximum Resolution	4096 x 2160 (4K)
OpenGL and Vulkan	In-situ Graphics only
CUDA and OpenCL Support	Supported

Features	Supported in NVIDIA AI Enterprise
ECC and Page Retirement	Supported
MIG GPU Support	Supported
Multi-vGPU	Supported
NVIDIA GPUDirect	Supported
Peer-to-Peer over NVLink	Supported
GPU Pass Through Support	Supported
Baremetal Support	Supported
AI and Data Science applications and Frameworks	Supported
Cloud Native ready	Supported

Note: Maximum 10 concurrent VMs per product license

The following table lists the ordering part numbers and feature codes.

Table 11. NVIDIA AI Enterprise Software (NVAIE)

Part number	Feature code 7S02CTO1WW	NVIDIA part number	Description
AI Enterprise Perpetual License			
7S02001BWW	S6YY	731-AI7004+P3CMI60	NVIDIA AI Enterprise Perpetual License and Support per GPU Socket, 5 Years
7S02001EWW	S6Z1	731-AI7004+P3EDI60	NVIDIA AI Enterprise Perpetual License and Support per GPU Socket, EDU, 5 Years
AI Enterprise Subscription License			
7S02001FWW	S6Z2	731-AI7003+P3CMI12	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, 1 Year
7S02001GWW	S6Z3	731-AI7003+P3CMI36	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, 3 Years
7S02001HWW	S6Z4	731-AI7003+P3CMI60	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, 5 Years
7S02001JWW	S6Z5	731-AI7003+P3EDI12	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, EDU, 1 Year
7S02001KWW	S6Z6	731-AI7003+P3EDI36	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, EDU, 3 Years
7S02001LWW	S6Z7	731-AI7003+P3EDI60	NVIDIA AI Enterprise Subscription License and Support per GPU Socket, EDU, 5 Years
AI Enterprise Support Services			
7S02001MWW	S6Z8	731-AI7007+P3CMI12	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, 1 Year
7S02001NWW	S6Z9	731-AI7007+P3CMI36	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, 3 Years
7S02001PWW	S6ZA	731-AI7007+P3CMI60	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, 5 Years
7S02001QWW	S6ZB	731-AI7007+P3EDI12	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, EDU, 1 Year
7S02001RWW	S6ZC	731-AI7007+P3EDI36	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, EDU, 3 Years

Part number	Feature code 7S02CTO1WW	NVIDIA part number	Description
-------------	----------------------------	--------------------	-------------

7S02001SWW	S6ZD	731-AI7007+P3EDI60	24X7 Support Services for NVIDIA Enterprise AI per GPU Socket, EDU, 5 Years
------------	------	--------------------	---

Find more information in the [NVIDIA AI Enterprise Sizing Guide](#).

NVIDIA Run:ai

NVIDIA Run:ai is an enterprise platform for AI workloads and GPU orchestration, accelerating AI operations with dynamic orchestration across the AI life cycle, maximizing GPU efficiency, and integrating seamlessly into hybrid AI infrastructure. The platform provides features such as AI-native workload orchestration, unified AI infrastructure management, flexible AI deployment, and open architecture, supporting public clouds, private clouds, hybrid environments, and on-premises data centers.

For more information, see the NVIDIA Run:ai product page:

<https://www.nvidia.com/en-us/software/run-ai/>

Table 12. NVIDIA Run:ai

Part number	Feature 7S02CTO1WW	NVIDIA part number	Description
Software subscription			
7S02004UWW	SDYT	744-RA7001+P3CMI12	NVIDIA Run:ai Subscription per GPU 1 Year
7S02004XWW	SDYW	744-RA7001+P3CMI36	NVIDIA Run:ai Subscription per GPU 3 Years
7S020050WW	SDYZ	744-RA7001+P3CMI60	NVIDIA Run:ai Subscription per GPU 5 Years
7S02004VWW	SDYU	744-RA7001+P3EDI12	NVIDIA Run:ai Subscription per GPU EDU 1 Year
7S02004YWW	SDYX	744-RA7001+P3EDI36	NVIDIA Run:ai Subscription per GPU EDU 3 Years
7S020051WW	SDZ0	744-RA7001+P3EDI60	NVIDIA Run:ai Subscription per GPU EDU 5 Years
7S02004WWW	SDYV	744-RA7001+P3INI12	NVIDIA Run:ai Subscription per GPU INC 1 Year
7S02004ZWW	SDYY	744-RA7001+P3INI36	NVIDIA Run:ai Subscription per GPU INC 3 Years
7S020052WW	SDZ1	744-RA7001+P3INI60	NVIDIA Run:ai Subscription per GPU INC 5 Years
Support Services subscription			
7S020053WW	SDZ2	744-RA7002+P3CMI12	24x7 Support Services for NVIDIA Run:ai Subscription per GPU 1 Year
7S020056WW	SDZ5	744-RA7002+P3CMI36	24x7 Support Services for NVIDIA Run:ai Subscription per GPU 3 Years
7S020059WW	SDZ8	744-RA7002+P3CMI60	24x7 Support Services for NVIDIA Run:ai Subscription per GPU 5 Years
7S020054WW	SDZ3	744-RA7002+P3EDI12	24x7 Support Services for NVIDIA Run:ai Subscription per GPU EDU 1 Year
7S02005AWW	SDZ9	744-RA7002+P3EDI60	24x7 Support Services for NVIDIA Run:ai Subscription per GPU EDU 5 Years
7S020057WW	SDZ6	744-RA7002+P3EDI36	24x7 Support Services for NVIDIA Run:ai Subscription per GPU EDU 3 Years
7S020055WW	SDZ4	744-RA7002+P3INI12	24x7 Support Services for NVIDIA Run:ai Subscription per GPU INC 1 Year
7S020058WW	SDZ7	744-RA7002+P3INI36	24x7 Support Services for NVIDIA Run:ai Subscription per GPU INC 3 Years
7S02005BWW	SDZA	744-RA7002+P3INI60	24x7 Support Services for NVIDIA Run:ai Subscription per GPU INC 5 Years

NVIDIA HPC Compiler Software

Table 13. NVIDIA HPC Compiler

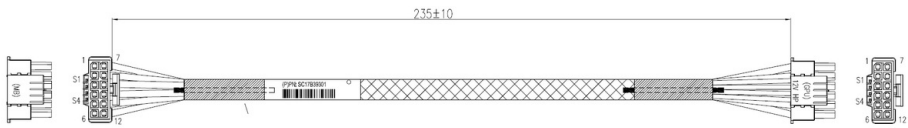
Part number	Feature code 7S09CTO6WW	Description
HPC Compiler Support Services		
7S090014WW	S924	NVIDIA HPC Compiler Support Services, 1 Year
7S090015WW	S925	NVIDIA HPC Compiler Support Services, 3 Years
7S09002GWW	S9UQ	NVIDIA HPC Compiler Support Services, 5 Years
7S090016WW	S926	NVIDIA HPC Compiler Support Services, EDU, 1 Year
7S090017WW	S927	NVIDIA HPC Compiler Support Services, EDU, 3 Years

Part number	Feature code 7S09CTO6WW	Description
7S09002HWW	S9UR	NVIDIA HPC Compiler Support Services, EDU, 5 Years
7S090018WW	S928	NVIDIA HPC Compiler Support Services - Additional Contact, 1 Year
7S09002JWW	S9US	NVIDIA HPC Compiler Support Services - Additional Contact, 3 Years
7S09002KWW	S9UT	NVIDIA HPC Compiler Support Services - Additional Contact, 5 Years
7S090019WW	S929	NVIDIA HPC Compiler Support Services - Additional Contact, EDU, 1 Year
7S09002LWW	S9UU	NVIDIA HPC Compiler Support Services - Additional Contact, EDU, 3 Years
7S09002MWW	S9UV	NVIDIA HPC Compiler Support Services - Additional Contact, EDU, 5 Years
HPC Compiler Premier Support Services		
7S09001AWW	S92A	NVIDIA HPC Compiler Premier Support Services, 1 Year
7S09002NWW	S9UW	NVIDIA HPC Compiler Premier Support Services, 3 Years
7S09002PWW	S9UX	NVIDIA HPC Compiler Premier Support Services, 5 Years
7S09001BWW	S92B	NVIDIA HPC Compiler Premier Support Services, EDU, 1 Year
7S09002QWW	S9UY	NVIDIA HPC Compiler Premier Support Services, EDU, 3 Years
7S09002RWW	S9UZ	NVIDIA HPC Compiler Premier Support Services, EDU, 5 Years
7S09001CWW	S92C	NVIDIA HPC Compiler Premier Support Services - Additional Contact, 1 Year
7S09002SWW	S9V0	NVIDIA HPC Compiler Premier Support Services - Additional Contact, 3 Years
7S09002TWW	S9V1	NVIDIA HPC Compiler Premier Support Services - Additional Contact, 5 Years
7S09001DWW	S92D	NVIDIA HPC Compiler Premier Support Services - Additional Contact, EDU, 1 Year
7S09002UWW	S9V2	NVIDIA HPC Compiler Premier Support Services - Additional Contact, EDU, 3 Years
7S09002VWW	S9V3	NVIDIA HPC Compiler Premier Support Services - Additional Contact, EDU, 5 Years

Auxiliary power cables

The RTX PRO 6000 Blackwell GPU option part number does not ship with auxiliary power cables. Cables are server-specific due to length requirements. For CTO orders, auxiliary power cables are derived by the configurator. For field upgrades, cables will need to be ordered separately as listed in the table below.

Table 14. Auxiliary power cables for RTX PRO 6000 Blackwell GPU

Auxiliary power cable needed with the SR675 V3	
235mm 16-pin (2x6+4) cable Option: 4X97A84510, ThinkSystem SR675 V3 Supplemental Power Cable for H100 GPU Option Feature: BSD2 SBB: SBB7A65299 Base: SC17B39301 FRU: 03LE554	

Regulatory approvals

The RTX PRO 6000 Blackwell GPU has the following regulatory approvals:

- RCM
- BSMI
- CE
- FCC
- ICES
- KCC
- cUL, UL
- VCCI

Operating environment

The RTX PRO 6000 Blackwell GPU has the following operating characteristics:

- Ambient temperature
 - Operational: 0°C to 50°C (-5°C to 55°C for short term*)
 - Storage: -40°C to 75°C
- Relative humidity:
 - Operational: 5-85% (5-93% short term*)
 - Storage: 5-95%

* A period not more than 96 hours consecutive, not to exceed 15 days per year.

Warranty

One year limited warranty. When installed in a Lenovo server, the GPU assumes the server's base warranty and any warranty upgrades.

Related publications

For more information, refer to these documents:

- ThinkSystem and ThinkAgile GPU Summary:
<https://lenovopress.lenovo.com/lp0768-thinksystem-thinkagile-gpu-summary>
- ServerProven compatibility:
<https://serverproven.lenovo.com/>
- RTX PRO 6000 Blackwell GPU product page:
<https://www.nvidia.com/en-us/data-center/rtx-pro-6000-blackwell-server-edition/>

Related product families

Product families related to this document are the following:

- [GPU adapters](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2025. All rights reserved.

This document, LP2263, was created or updated on September 23, 2025.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP2263>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP2263>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ServerProven®

ThinkAgile®

ThinkSystem®

The following terms are trademarks of other companies:

AMD is a trademark of Advanced Micro Devices, Inc.

Intel® is a trademark of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

DirectX®, Windows Server®, and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.